A large-scale analysis of racial disparities in police stops across the United States

Emma Pierson
Cornell Tech
@2plus2make5

New York Statewide Civil Legal Aid Technology Conference
April 11, 2022

Police traffic stops

- **20M** Americans stopped every year
  - One of the most common ways we interact with police

[1] Baumgartner, Epp, and Shoub. "Suspect citizens: What 20 million traffic stops tell us about policing and race". Cambridge University Press, 2018.
[2] Davis, Whyde, and Langton. "Contacts between police and the public, 2015". Technical report, Bureau of Justice Statistics, 2018.

Are traffic stops racially discriminatory?

- Concerns that traffic stops are **racially discriminatory**
- Racial discrimination = when someone is treated more negatively because of their race
- Difficult to statistically test for

[1] Epp, Maynard-Moody, and Haider-Markel. “Pulled over: how police stops define race and citizenship”. University of Chicago Press, 2014.

Creating a national traffic stop dataset

- Unified dataset didn’t exist!

Creating a national traffic stop dataset

- **1. Obtaining data**: journalists submit data requests to more than **150** police departments over **5 years**



Creating a national traffic stop dataset

- **2. Standardizing data:** many data formats, coding conventions, errors -> **thousands of hours** of cleanup

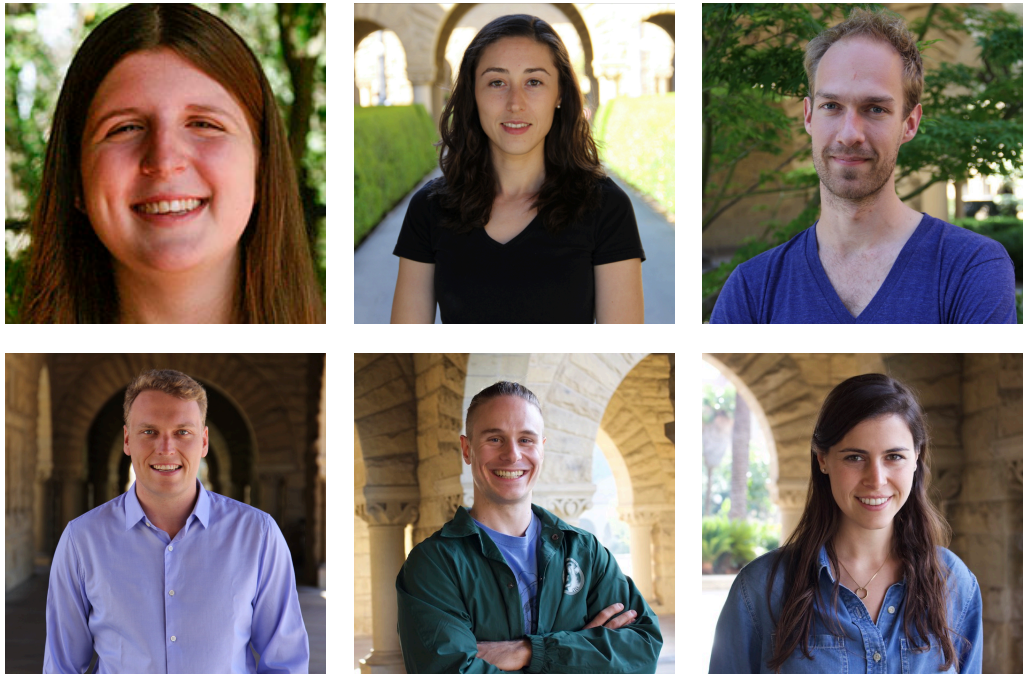

Creating a national traffic stop dataset

- All data and code available: openpolicing.stanford.edu

	Stops	City agencies	State agencies
Full dataset	255M	56	33
Main analysis	95M	35	21

Analysis

- 1. Are police **stops** discriminatory?
- 2. Are police **searches** discriminatory?
- 3. Can **policy changes** (marijuana legalization) affect racial disparities?

Analysis

- 1. Are police **stops** discriminatory?
- 2. Are police **searches** discriminatory?
- 3. Can **policy changes** (marijuana legalization) affect racial disparities?

Discrimination in police traffic stops

- Are police racially profiling when they decide who to stop?
- Just comparing stops per capita is not really enough because some races might drive more, commit more violations, etc
- **Idea**: use the fact that it's harder to tell someone's race when it's dark -> harder to discriminate on basis of race

Grogger and Ridgeway. "Testing for racial profiling in traffic stops from behind a veil of darkness". *JASA*, 2006.
Pierson et al. "A large scale analysis of racial disparities in police stops across the United States". *Nature Human Behavior*, 2020.

Discrimination in police traffic stops

Clock time the same; darkness different

6 PM June 14

6 PM November 14

Discrimination in police traffic stops

Clock time the same; darkness different

Worried about seasonal effects? Just use data near daylight savings time to get **natural experiment**

6 PM November 2

6 PM November 3

Discrimination in police traffic stops

$$Pr(\text{stopped driver is black}) = \text{logistic}(\beta_{\text{darkness}} d + controls)$$

Controls: clock time; location; Spring/Fall DST

Fit single model for state + city stops, but estimate separate darkness effects

6 PM November 2

6 PM November 3

Discrimination in police traffic stops

$$Pr(\text{stopped driver is black}) = \text{logistic}(\beta_{\text{darkness}} d + controls)$$

β_{darkness} is **−0.033** (−0.039, −0.027) for state patrol stops,
−0.039 (−0.045, −0.033) for city stops

6 PM November 2

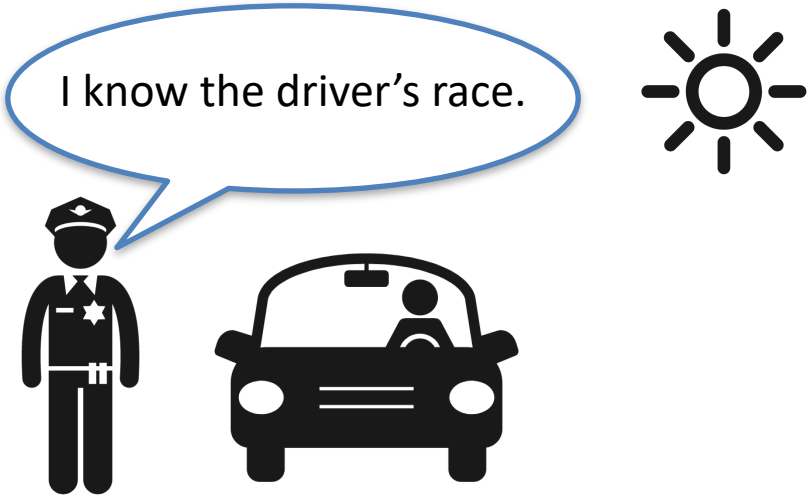

6 PM November 3

Discrimination in police traffic stops

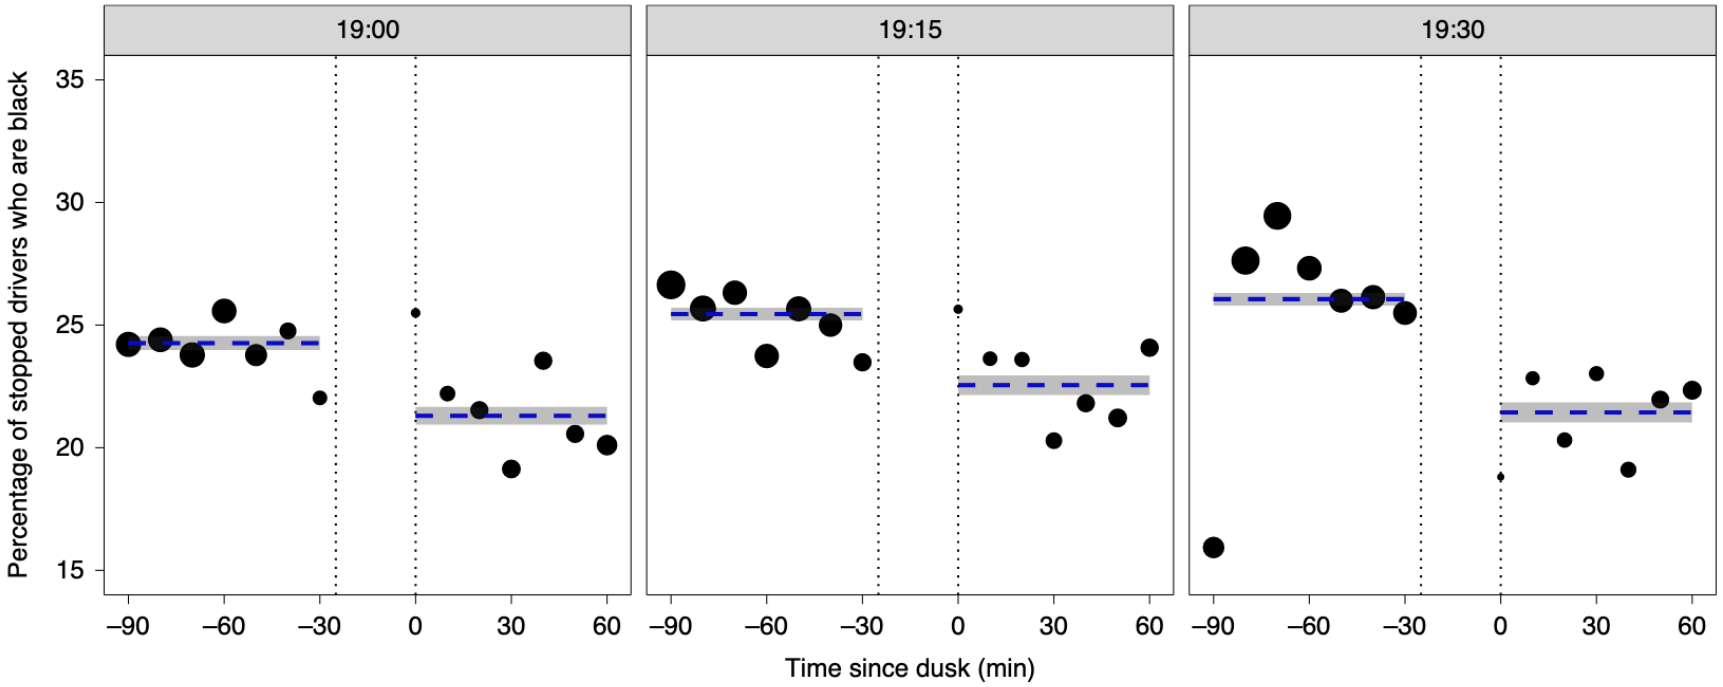

Illustration of core idea with data from Texas

Analysis

- 1. Are police **stops** discriminatory?
- 2. Are police **searches** discriminatory?
- 3. Can **policy changes** (marijuana legalization) affect racial disparities?

Are police searches discriminatory?

- After stopping a driver, police may conduct a search to **find contraband** (drugs, weapons…)
- We test whether minorities are searched when **less likely to have contraband** (lower threshold of evidence)
- Not trying to test for all forms of problematic police behavior

Simple test I: look at search rates

- Minority drivers are more likely to be searched after stop than white drivers

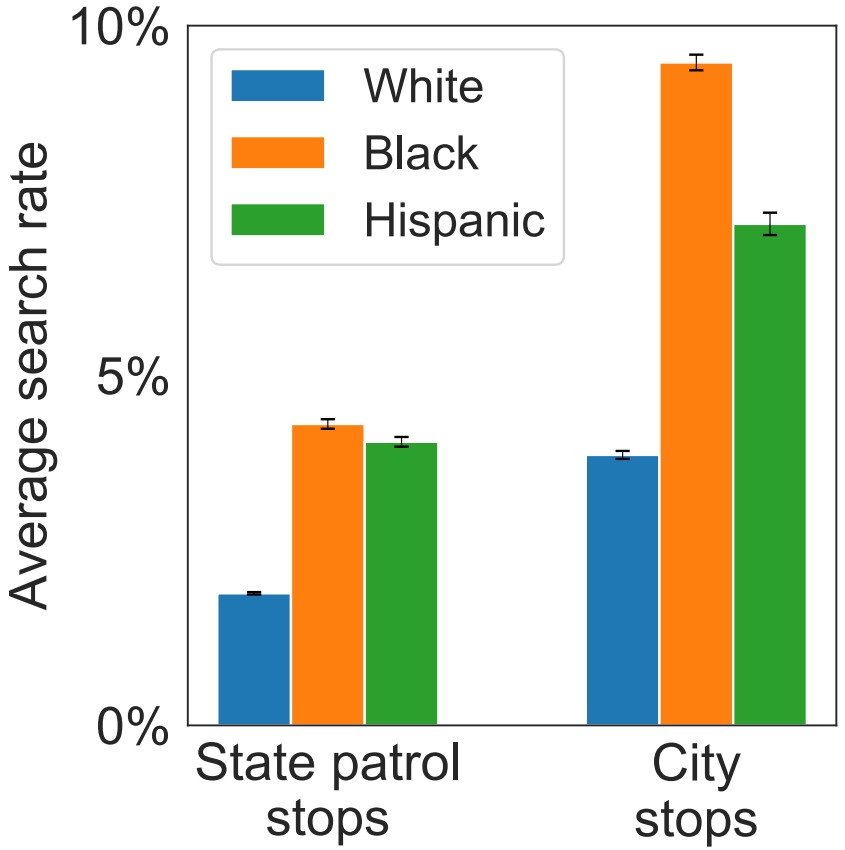

Simple test I: look at search rates

- Minority drivers are more likely to be searched after stop than white drivers
- **Flaw:** higher search rates don't prove police are being discriminatory
- Some races may be more likely to carry contraband (weapons, drugs…)

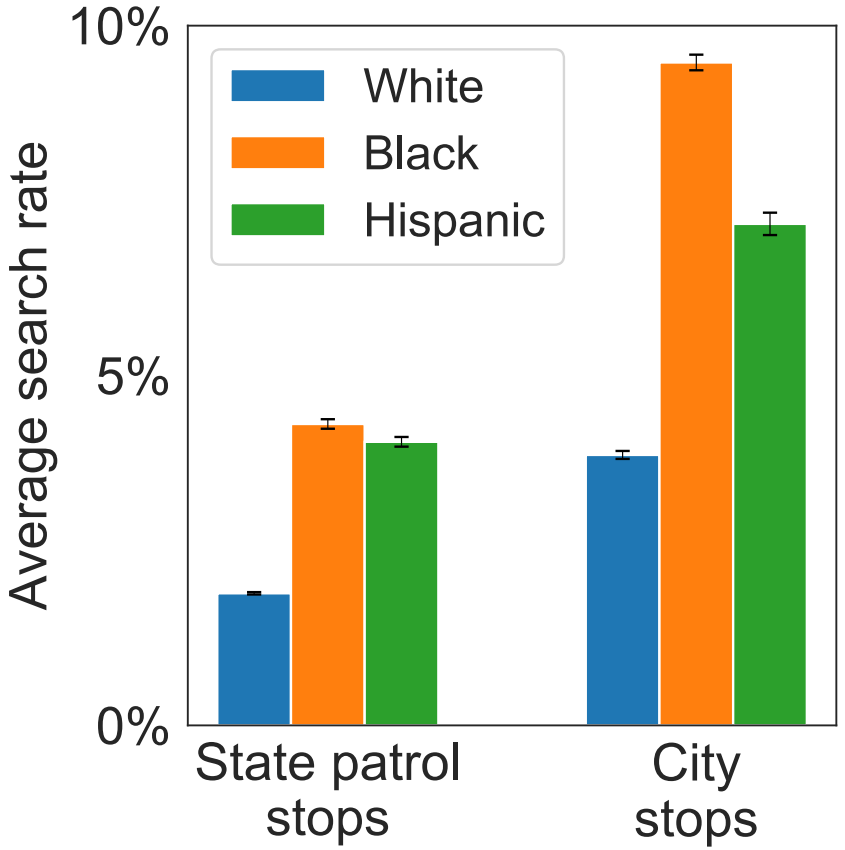

Simple test II: look at search outcomes

- **Outcome test**: look at how likely searches are to find contraband (**hit rate**)
- Hit rates vary by race = discrimination

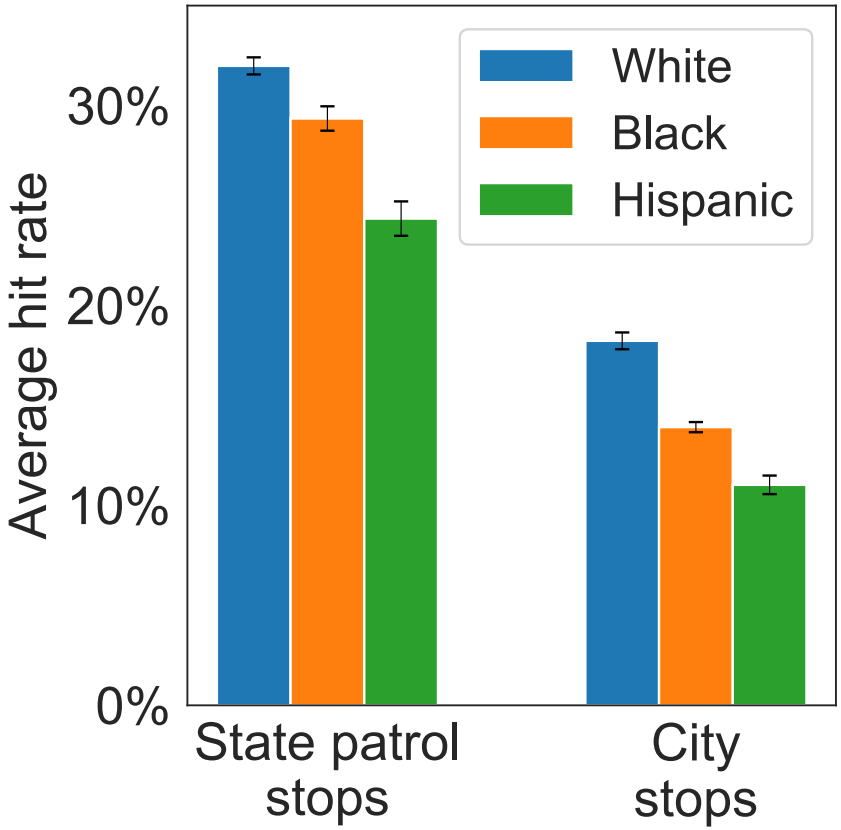

Gary Becker. *The Economics of Discrimination*. University of Chicago Press, 1957.

Flaw: infra-marginality

Hypothetical example:

	Likely to carry contraband	**Unlikely to carry contraband**
Black drivers	50% carry	5% carry
White drivers	75% carry	5% carry

Police search everyone >**10%** likely to carry contraband

What are hit rates for white and black drivers?

[1] Ayres. "Outcome tests of racial disparities in police practices". *Justice Research and Policy*, 2002.
[2] Anwar and Fang. "An alternative test of racial prejudice in motor vehicle searches: Theory and evidence". *The American Economic Review*, 2006.

Flaw: infra-marginality

Hypothetical example:

	Likely to carry contraband	**Unlikely to carry contraband**
Black drivers	50% carry	5% carry
White drivers	75% carry	5% carry

Police search everyone >**10%** likely to carry contraband

What are hit rates for white and black drivers?

Solution: use Bayesian latent variable model to infer threshold

[1] Ayres. "Outcome tests of racial disparities in police practices". *Justice Research and Policy*, 2002.
[2] Anwar and Fang. "An alternative test of racial prejudice in motor vehicle searches: Theory and evidence". *The American Economic Review*, 2006.

More complicated: threshold test

- Stylized model of police stop
- **Goal:** estimate search thresholds that are consistent with observed data (search and hit rates)
- Discrimination = lower search thresholds for minority drivers

Simoiu, Corbett-Davies, and Goel. “The problem of infra-marginality in outcome tests for discrimination”. *Annals of Applied Statistics*, 2017.

Threshold test model of a police stop

- 1. Officer estimates probability *p* each stopped person carries contraband. *p* drawn from **risk distribution**

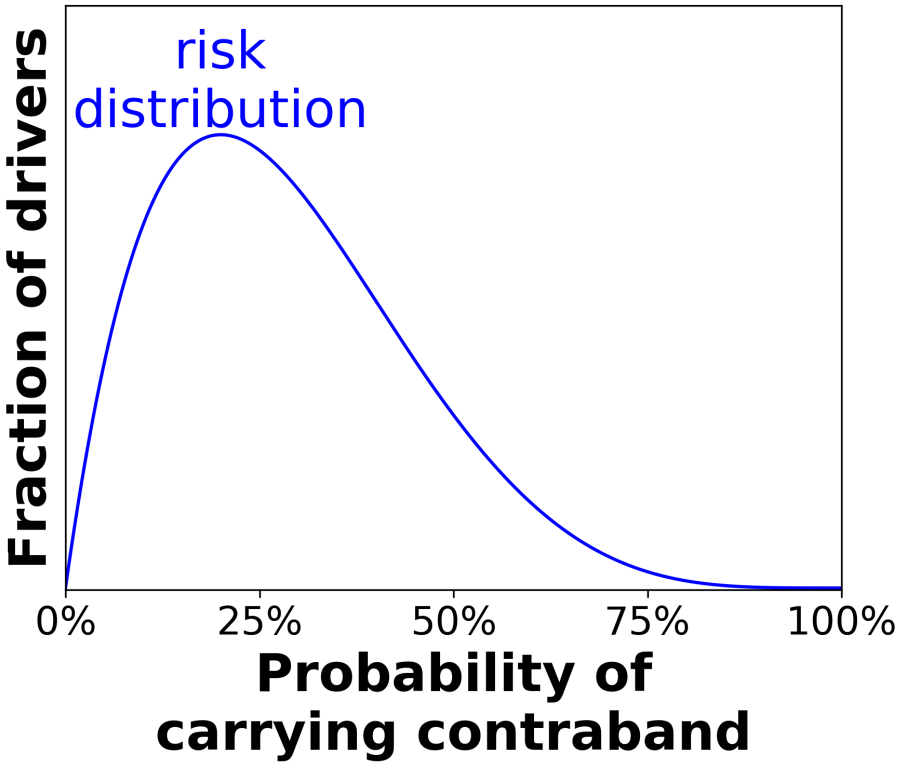

Threshold test model of a police stop

- 1. Officer estimates probability *p* each stopped person carries contraband. *p* drawn from **risk distribution**

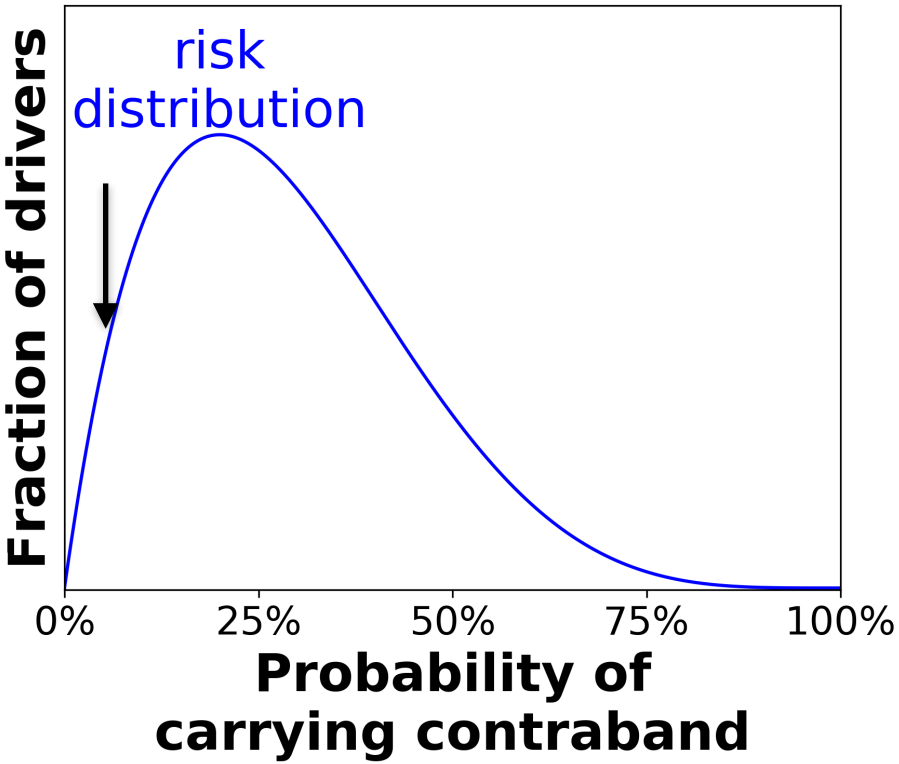

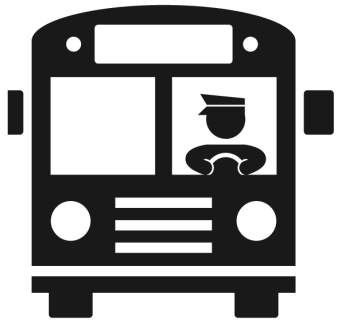

Threshold test model of a police stop

- 1. Officer estimates probability *p* each stopped person carries contraband. *p* drawn from **risk distribution**

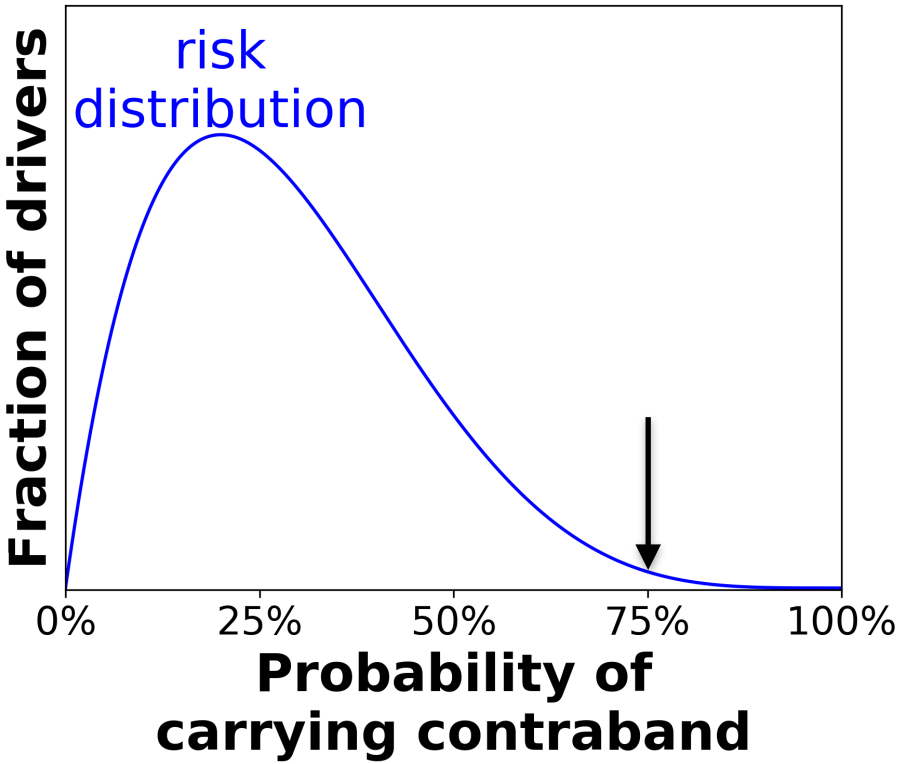

Threshold test model of a police stop

- 1. Officer estimates probability *p* each stopped person carries contraband. *p* drawn from **risk distribution**

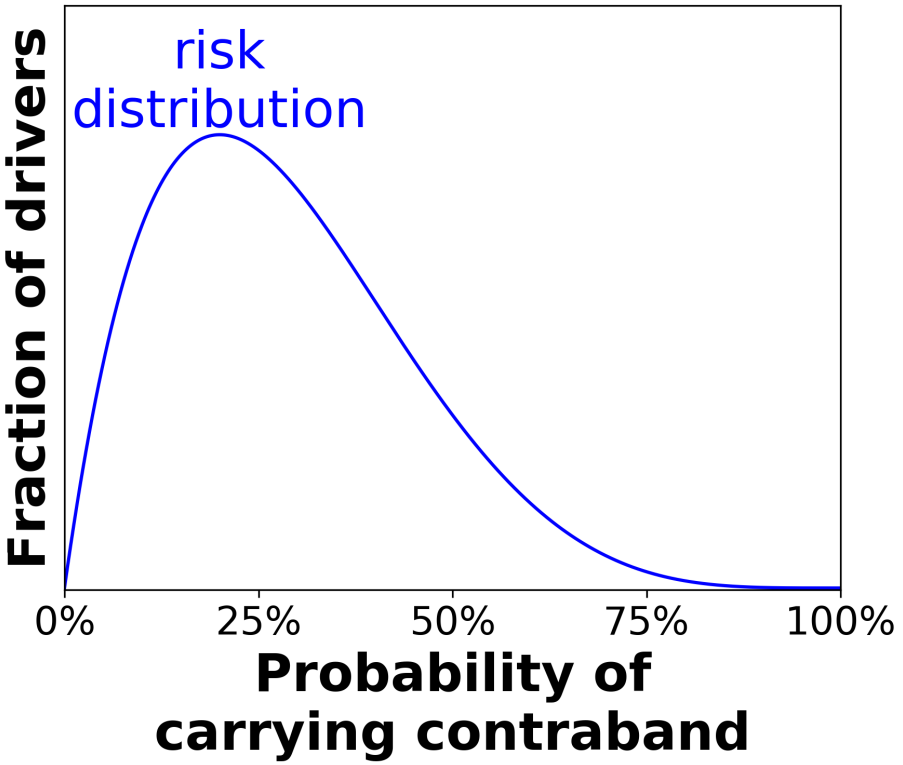

Threshold test model of a police stop

- 1. Officer estimates probability *p* each stopped person carries contraband. *p* drawn from **risk distribution**
- 2. If *p* exceeds a **threshold,** officer searches person and finds contraband with probability *p*

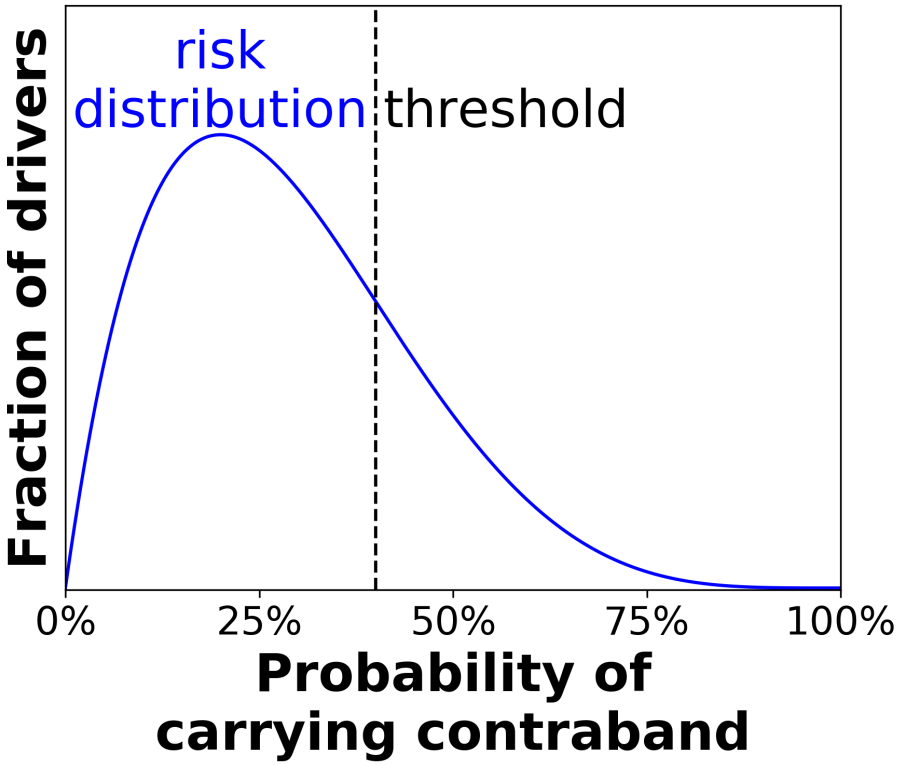

Threshold test model of a police stop

- 1. Officer estimates probability *p* each stopped person carries contraband. *p* drawn from **risk distribution**
- 2. If *p* exceeds a **threshold,** officer searches person and finds contraband with probability *p*

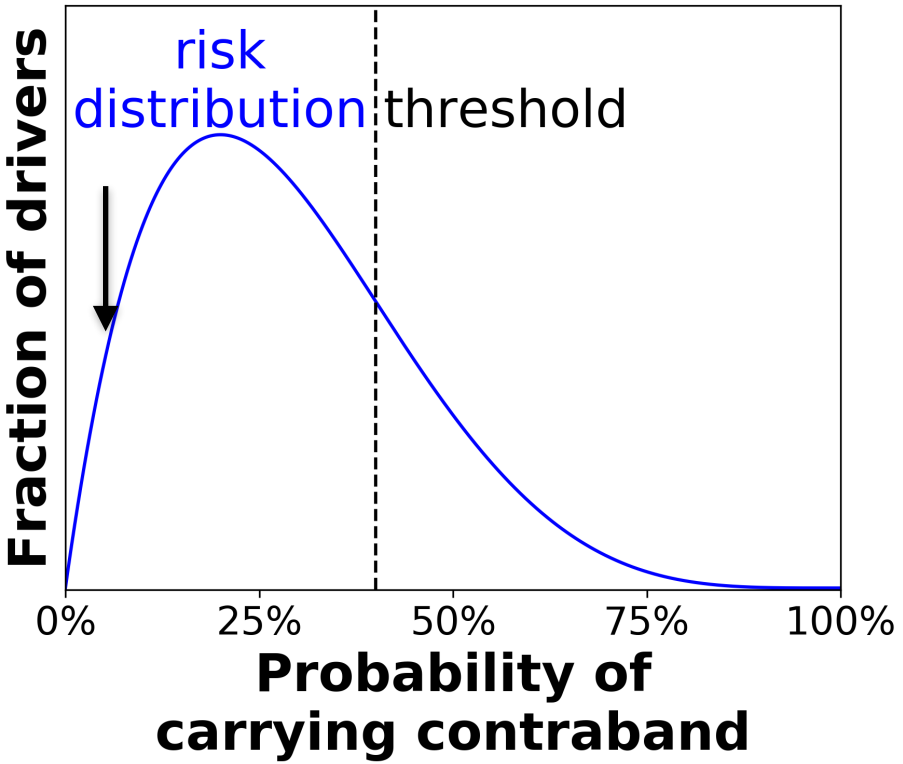

Below threshold: no search

Threshold test model of a police stop

- 1. Officer estimates probability *p* each stopped person carries contraband. *p* drawn from **risk distribution**
- 2. If *p* exceeds a **threshold,** officer searches person and finds contraband with probability *p*

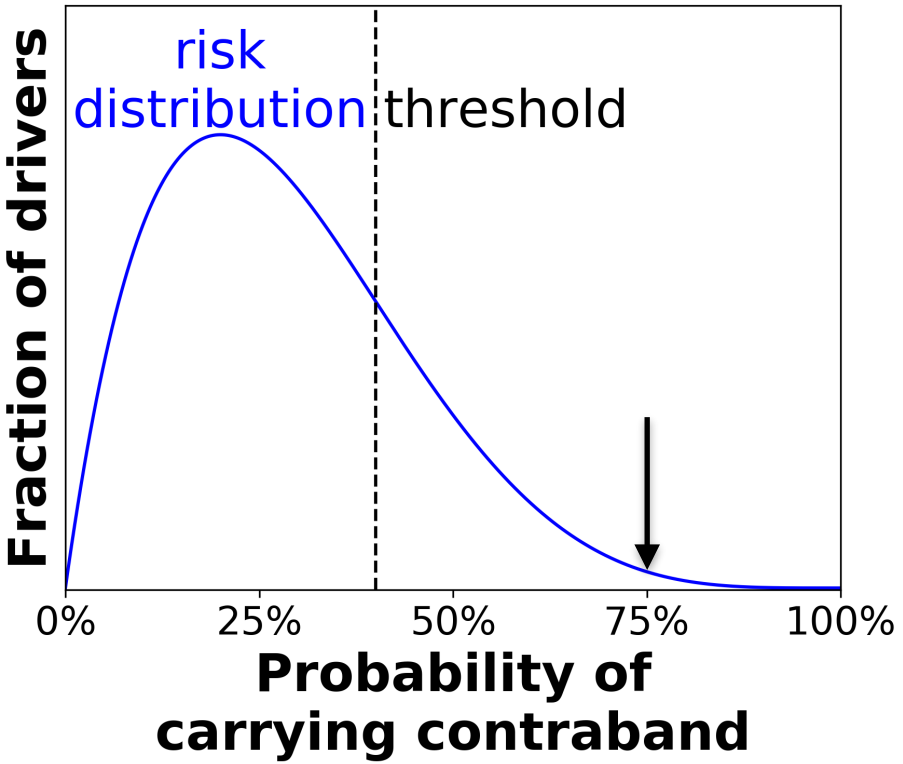

Above threshold:
Search,
and 75% chance of finding contraband

Threshold test model of a police stop

- 1. Officer estimates probability *p* each stopped person carries contraband. *p* drawn from **risk distribution**
- 2. If *p* exceeds a **threshold,** officer searches person and finds contraband with probability *p*

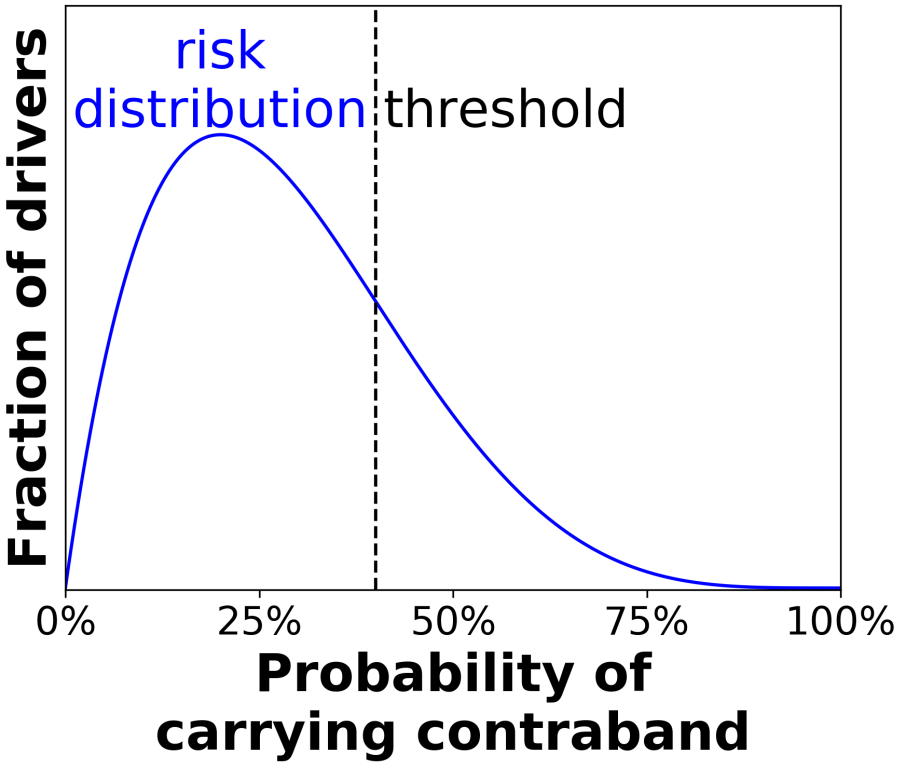

Thresholds + risk distributions vary by race and location

Discrimination = lower thresholds for minority drivers

Problem: model fitting is **slow**!

- Describe a **new family of probability distributions** (“discriminant” distributions)
- Use them to create **fast threshold test** (**100x** faster)
  - Allows us to scale to national dataset
  - Allows non-academic practitioners to use test
- A story for another time…
- **Takeaway**: probability distributions matter!

Pierson, Corbett-Davies, and Goel. “Fast Threshold Tests for Detecting Discrimination”. AISTATS, 2018 (**best paper award**).

With fast threshold test, we can fit a national model!

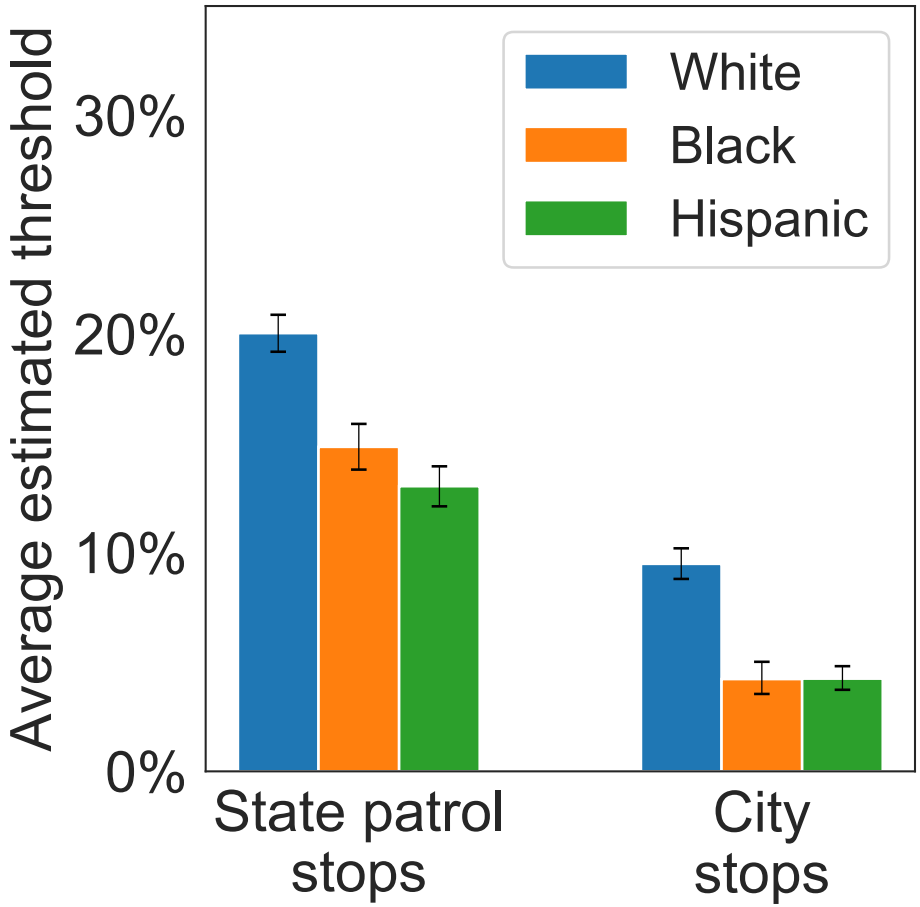

Thresholds are averaged across locations.

Posterior predictive checks used to ensure model fits the data

Summary of search analysis

Search rates are **higher** for minorities…

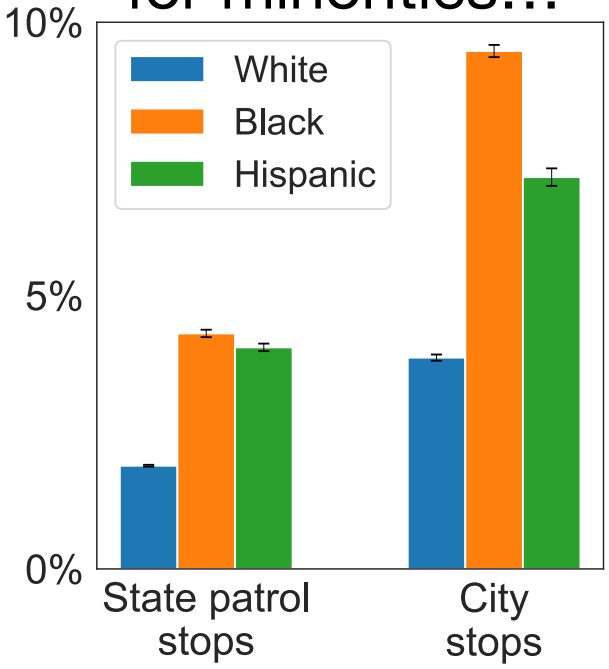

Hit rates are **lower…**

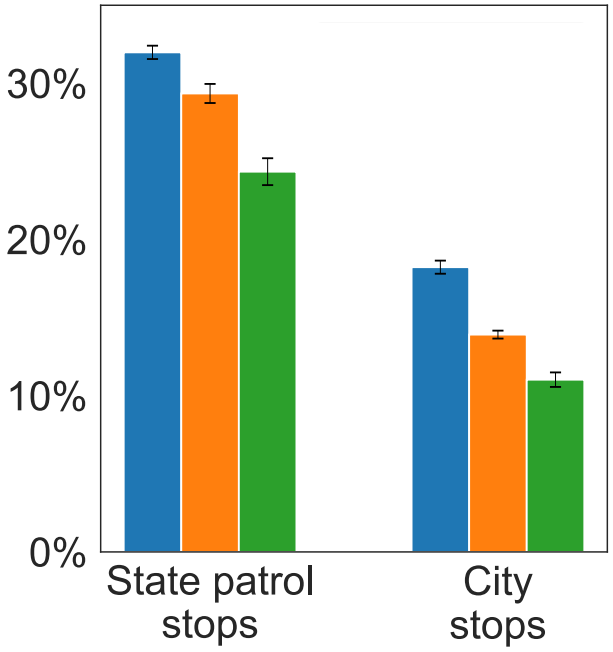

Thresholds are **lower**

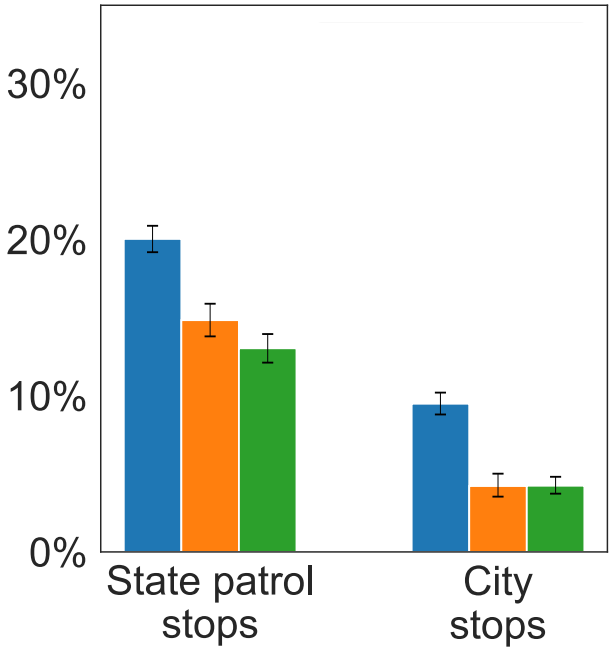

A **characteristic pattern** for discriminatory searches.

All three tests suggest discrimination against minorities

but threshold test is **robust to statistical flaws of simpler tests**

Summary of search analysis

Search rates are **higher** for minorities…

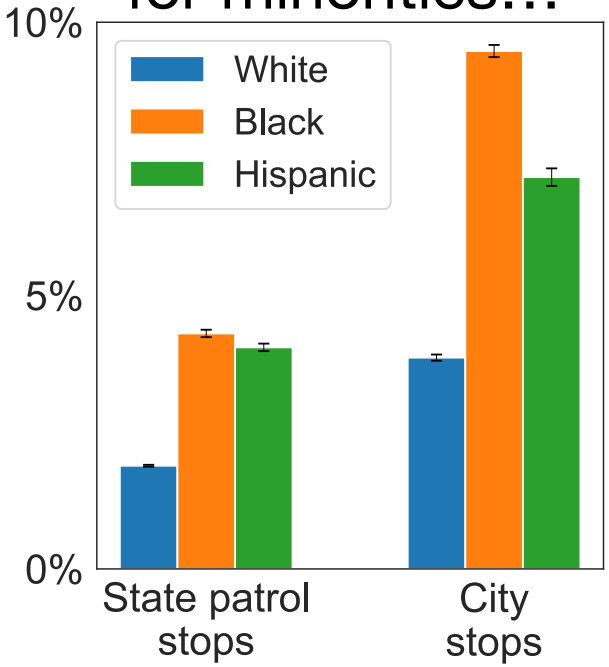

Hit rates are **lower…**

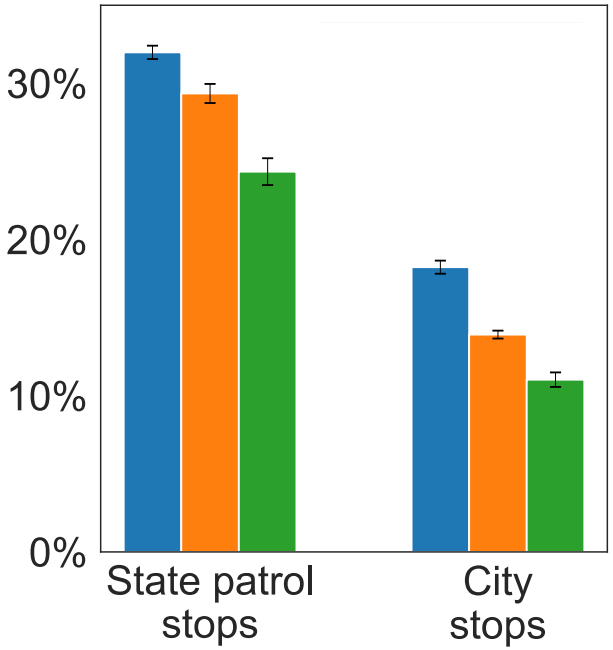

Thresholds are **lower**

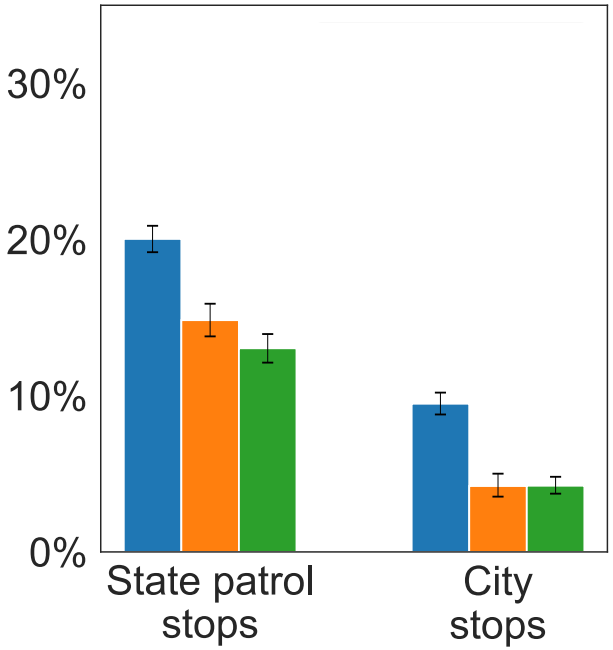

Same methods apply in other datasets with decisions and outcomes

Medical tests: Decision: gets tested? Outcome: tests positive?

Loans: Decision: receives loan? Outcome: repays loan?

Analysis

- 1. Are police **stops** discriminatory?
- 2. Are police **searches** discriminatory?
- 3. Can **policy changes** (marijuana legalization) affect racial disparities?

Difference-in-differences

- What is the effect of legalization of marijuana on whether drivers are searched after a stop?
- Compare change in search rates in two states where marijuana was legalized to 12 states where it was not

Difference-in-differences

"Treatment" states where marijuana was legalized

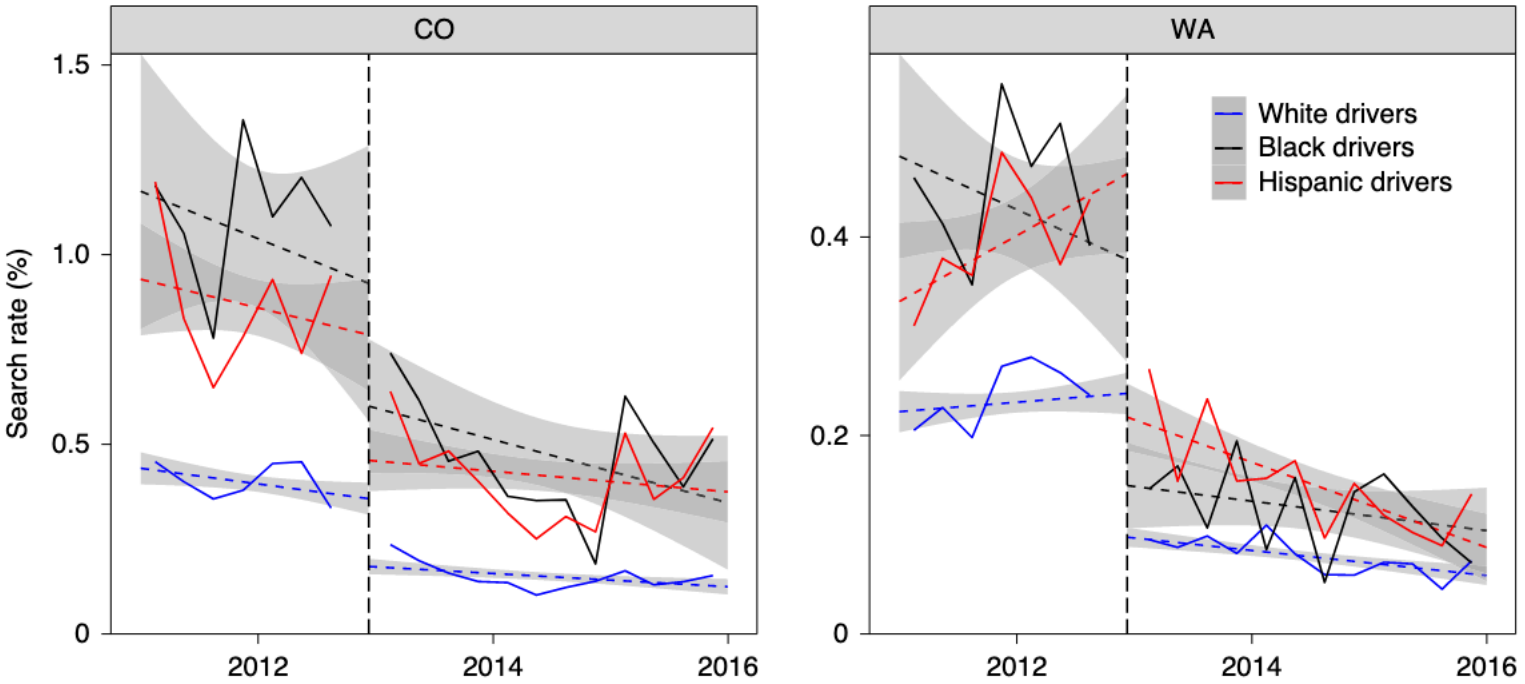

Difference-in-differences

"Control" states where marijuana was not legalized

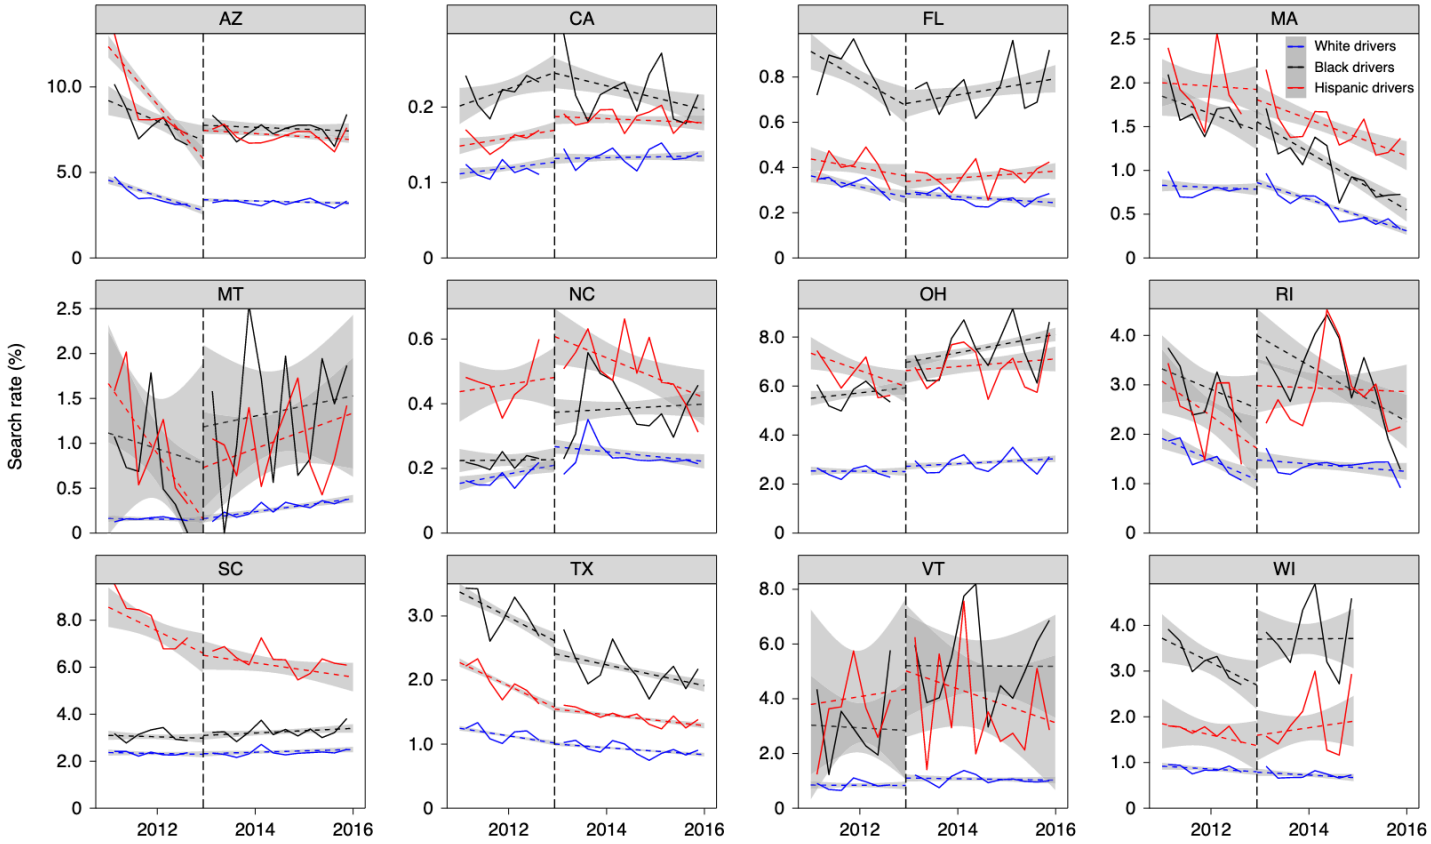

Difference-in-differences

Table 1 | Effects of legalization of recreational marijuana on search rates, as estimated using a difference-in-difference model. All race groups experienced a large drop in search rate

	Coefficient	s.e.	95% CI	*P* value
Effect of legalization on white drivers	-0.96	0.02	(-1.01, -0.92)	<0.001
Effect of legalization on black drivers	-0.98	0.06	(-1.09, -0.87)	<0.001
Effect of legalization on Hispanic drivers	-0.69	0.03	(-0.76, -0.63)	<0.001
Time (years)	-0.04	0.00	(-0.04, -0.04)	<0.001
Black driver	0.78	0.00	(0.77, 0.79)	<0.001
Hispanic driver	0.59	0.00	(0.58, 0.59)	<0.001

Public policy impact

October 8, 2019

Los Angeles Times LOG IN

CALIFORNIA

LAPD searches blacks and Latinos more. But they're less likely to have contraband than whites

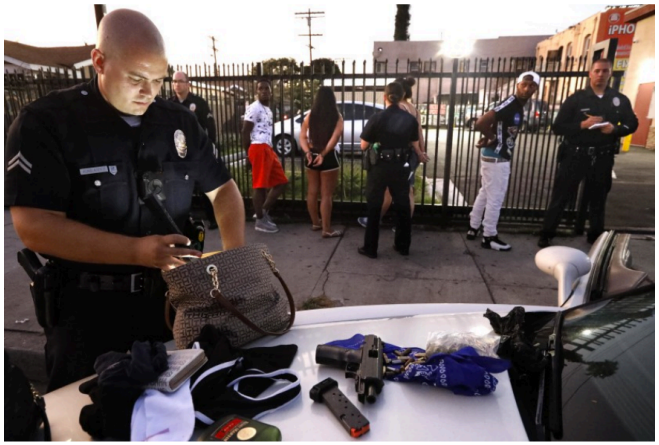

Public policy impact

October 8, 2019

Los Angeles Times LOG IN

CALIFORNIA

LAPD searches blacks and Latinos more. But they're less likely to have contraband than whites

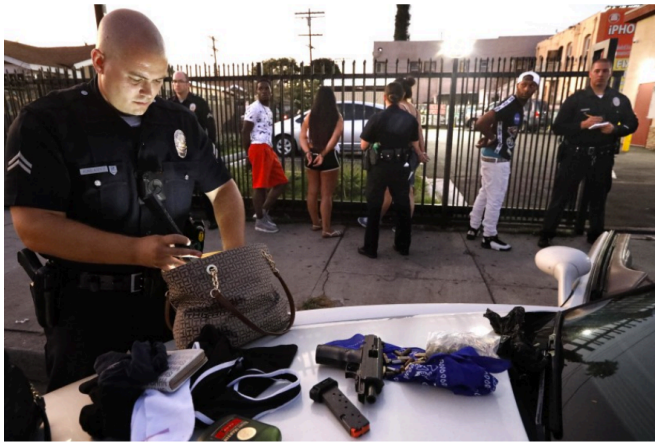

October 13, 2019

Los Angeles Times LOG IN

CALIFORNIA

LAPD will drastically cut back on pulling over random vehicles over racial bias concerns

Thank you!

Do you have datasets which might benefit from similar statistical analyses?
Please reach out.
emma.pierson@cornell.edu